

# Perceus – A Cluster Provisioning and Management Toolkit

Bill Strossman



# What Perceus Is

- An open source cluster toolkit for provisioning and managing stateless or stateful clusters
- Highly scalable
- Highly flexible and customizable
- Two software packages (one for the head node, one for the workers)
- Comprehensive set of command-line utilities.
- Easy to set up and use

# What Perceus Is Not

- A complete Linux distribution – however, there are two Linux distributions that include perceus
  - Caos NSA - this is a Redhat variant, so software that runs on RHEL/CentOS/SL should run on Caos NSA
  - GravityOS – this is debian based
- A GUI toolkit – however, a web GUI is available beginning with perceus 1.5.
- A cluster monitoring tool

# Supported Operating Systems

In addition to Caos NSA and GravityOS perceus is officially supported on:

- RHEL 4.5 (and newer) and RHEL 5.x
- Debian Squeeze

Also known to work on:

- CentOS and Scientific Linux
- Ubuntu
- SuSe
- possibly others

# Definitions

- Node - compute node
- VNFS (Virtual Node File System) - the template image used to provision nodes.
- VNFS capsule – packaged-up version of VNFS; contains all information necessary for provisioning.
- VNFS image – actual filesystem that is transported to and installed on each node.

# Definitions - continued

- Stateless – the operating system is never installed to the node's local media (if any), so the node must be provisioned on each boot. However, a stateless OS can make use of the node's local disk partitions for scratch, spooling, Hadoop data, etc.
- Hybridization – some of the OS is NFS-exported to the nodes (/usr/share is a frequent example)

# VNFS Capsules

- VNFS capsules for GravityOS, Caos NSA, and CentOS 5 are downloadable from [www.perceus.org](http://www.perceus.org). However, perceus contains easy-to-use, customizable shell scripts for building your own.
- Different nodes can be provisioned with different VNFS capsules. Our T3 cluster uses two VNFS capsules – one with hadoop, the other without.
- The VNFS capsule OS does not need to be the same as the head node OS.

# Installation

Install required perl packages

- perl-IO-Interface-0.98-1.redhat.noarch.rpm
- perl-Net-ARP-1.0-2.redhat.noarch.rpm
- perl-Unix-Syslog-0.99-1.redhat.noarch.rpm

Install perceus on head node

- perceus-1.5.2-2111.x86\_64.rpm

If you want the web GUI, install it now

- perceus-cgi-1.5.2-2111.x86\_64.rpm

Download link:



# Configuration

Edit the file `/etc/perceus/perceus.conf`. You only need to change two lines

- `master network device = virbr0`
- `vnfs transfer master = 192.168.122.1`

The defaults for everything else should work.

Now you need to create a VNFS

- `cd /usr/share/perceus/vnfs-scripts`
- `./centos-5.1-genchroot.sh`
- `chroot2stateless.sh \`  
`/var/tmp/vnfs/centos-5.1-1.x86_64 \`  
`/root/centos-5.1-1.stateless.x86_64.vnfs`

# Configuration

Now you need to initialize perceus

- `perceus init`
- Answer the questions

Import the VNFS

- `perceus vnfs import \`  
`/root/centos-5.1-1.stateless.x86_64.vnfs`

Mount the VNFS

- `perceus vnfs mount \`  
`centos-5.1-1.stateless.x86_64`

# Configuration

Now you can modify the VNFS. It is located at  
`/mnt/centos-5.1-1.stateless.x86_64`

- `rpm -ivh \`  
`--root /mnt/centos-5.1-1.stateless.x86_64 \`  
`perceus-provisiond-1.5.2-2111.x86_64.rpm`
- `yum update \`  
`--installroot=/mnt/centos-5.1-1.stateless.x86_64`

Time for a few important steps

- `yum clean all \`  
`--installroot=/mnt/centos-5.1-1.stateless.x86_64`  
`rpm -ev --root /mnt/centos-5.1-1.stateless.x86_64 \`  
`<old kernels>`

# Configuration

You can now perform any additional customization on the VNFS

- `rm -rf /mnt/centos-5.1-1.stateless.x86_64/usr/share/man`
- `rm -rf /mnt/centos-5.1-1.stateless.x86_64/usr/share/doc`
- `rm -rf /mnt/centos-5.1-1.stateless.x86_64/usr/share/info`
- add mount points to fstab for local partitions, etc.

Unmount the VNFS

- `perceus umount centos-5.1-1.stateless.x86_64`

Restart perceus (just to be safe)

- `/etc/init.d/perceus restart`

# Booting Nodes

Now boot the nodes in the order you want them to be numbered (you only need to do this once)

On the head node you need to set the vnfs for the nodes

- `perceus node set vnfs centos-5.1-1.stateless.x86_64 node01`
- `perceus node set vnfs centos-5.1-1.stateless.x86_64 node02`
- for many nodes one can use `node00[00-09,10-40]`
- This step only needs to be done once as nodes are added

Done!

# Concluding Remarks

- Perceus has its own DHCP server
- As changes are made to the VNFS, the updated image can be “pushed out” to the nodes without rebooting using the perceus “livesync” utility (however, kernel updates require a reboot).
- Can be compiled from source with rpmbuild
- Can back up your VNFS with “perceus vnfs export” utility.
- Can have many different VNFS capsules for a heterogeneous cluster or subclusters.